

# Mining Rooted Ordered Trees under Subtree Homeomorphism

Mostafa Haghiri Chehreghani ·  
Maurice Bruynooghe

Received: date / Accepted: date

**Abstract** Mining frequent tree patterns has many applications in different areas such as XML data, bioinformatics and World Wide Web. The crucial step in frequent pattern mining is frequency counting, which involves a matching operator to find occurrences (instances) of a tree pattern in a given collection of trees. A widely used matching operator for tree-structured data is *subtree homeomorphism*, where an edge in the tree pattern is mapped onto an ancestor-descendant relationship in the given tree. Tree patterns that are frequent under subtree homeomorphism are usually called *embedded patterns*. In this paper, we present an efficient algorithm for subtree homeomorphism with application to frequent pattern mining. We propose a compact data-structure, called **occ**, which stores only information about the rightmost paths of occurrences and hence can encode and represent several occurrences of a tree pattern. We then define efficient join operations on the **occ** data-structure, which help us count occurrences of tree patterns according to occurrences of their proper subtrees. Based on the proposed subtree homeomorphism method, we develop an effective pattern mining algorithm, called TPMiner. We evaluate the efficiency of TPMiner on several real-world and synthetic datasets. Our extensive experiments confirm that TPMiner always outperforms well-known existing algorithms, and in several cases the improvement with respect to existing algorithms is significant.

**Keywords** XML data · rooted ordered trees · frequent tree patterns · subtree homeomorphism · embedded subtrees

## 1 Introduction

Many semi-structured data such as XML documents are represented by rooted ordered trees. One of the most important problems in the data mining of these tree-structured data is frequent tree pattern discovery. Mining frequent tree patterns is

---

Mostafa Haghiri Chehreghani · Maurice Bruynooghe  
Department of Computer Science, KU Leuven, 3001 Leuven, Belgium  
Mostafa.HaghiriChehreghani@cs.kuleuven.be, Maurice.Bruynooghe@cs.kuleuven.be

very useful in various domains such as network routing [11], bioinformatics [26] and user web log data analysis [14]. Furthermore, it is a crucial step in several other data mining and machine learning problems such as clustering and classification [27].

In general, algorithms proposed for finding frequent tree patterns include two main phases: 1) generating candidate tree patterns, and 2) counting the frequency of every generated tree pattern in a given collection of trees (called the *database trees* from now on). The generation step (which involves a refinement operator) is computationally easy. There are methods, such as *rightmost path extension*, that can generate efficiently all non-redundant rooted ordered trees, i.e., each in  $O(1)$  computational time [26] and [1]. The frequency counting step is computationally expensive. Empirical comparison of these two phases can be found e.g., in [10], where Chi et. al. showed that a significant part of the time required for finding frequent patterns is spent on frequency counting. Thereby, the particular method used for frequency counting can significantly affect the efficiency of the tree mining algorithm.

The frequency counting step involves a *matching operator* [15]. A widely used matching operator between a tree pattern and a database tree is *subtree homeomorphism*; an injective mapping of the vertices such that an edge in the tree pattern is mapped onto an ancestor-descendant relationship in the database tree. Frequent tree patterns under subtree homeomorphism are called *embedded patterns*. They have many applications in different areas. Zaki et. al. [27] presented XRules, a classifier based on frequent embedded tree patterns, and showed its high performance compared to classifiers such as SVM. Ivancsy and Vajk used frequent embedded tree patterns for analyzing the navigational behavior of the web users [14].

Two widely used frequency notions are *per-tree frequency*, where only the occurrence of a tree pattern inside a database tree is important; and *per-occurrence frequency*, where the number of occurrences is important, too. While there exist algorithms optimized for the first notion [26], [22] and [23], this notion is covered also by the algorithms proposed for the second notion. Per-occurrence frequency counting is computationally more expensive than per-tree frequency counting. In the current paper, our concern is *per-occurrence frequency*. An extensive discussion about the applications in which the second notion is preferred can be found e.g., in [20]. One of the investigated examples is a digital library where author information are separately stored in database trees in some form, e.g., *author–book–area–publisher*. A user may be interested in finding out information about the popular publishers of every *area*. Then, the repetition of items within a database tree becomes important, hence, per-occurrence frequency is more suitable than per-tree frequency [20].

Two categories of approaches have been used for counting occurrences of tree patterns under subtree homeomorphism. The first category includes approaches that use one of the algorithms proposed for subtree homeomorphism between two trees. HTreeMiner [26] employs such an approach. The second category includes approaches that store the information representing/encoding the occurrences of tree patterns. Then, when the tree pattern is extended to a larger one, its stored information is also extended, in a specific way, to represent the occurrences of the extended pattern. These approaches are sometimes called *vertical* approaches. VTreeMiner [26] and MB3Miner [20] are examples of the methods that use such vertical approaches. As studied in [26], vertical approaches are more efficient than the approaches in the first

category. Many efficient vertical algorithms are based on the numbering scheme proposed by Dietz [13]. This scheme uses a tree traversal order to determine the ancestor-descendant relationship between pairs of vertices. It associates each vertex with a pair of numbers, sometimes called *scope*. For instance, VTreeMiner and TreeMinerD [26] and TwigList [18] use this scheme in different forms, to design efficient methods for counting occurrences of tree patterns.

The problem with these algorithms is that in order to count all occurrences, they use data-structures that represent whole occurrences. This renders the algorithms inefficient, especially when patterns are large and have many occurrences in the database trees. In the worst case, the number of occurrences of a tree pattern can be exponential in terms of the size of pattern and database [9]. Therefore, keeping track of all occurrences can significantly reduce the efficiency of the algorithm, in particular when tree patterns have many occurrences in the database trees.

The main contribution of the current paper is to introduce a novel vertical algorithm for the class of rooted ordered trees. It uses a more compact data-structure, called **occ** (an abbreviation for **occurrence compressor**) for representing occurrences, and a more efficient subtree homeomorphism algorithm based on Dietz’s numbering scheme [13]. An **occ** data-structure stores only information about rightmost paths of occurrences and hence can represent/encode all occurrences that have the rightmost path in common. The number of such occurrences can be exponential, even though the size of the **occ** is only  $O(d)$ , where  $d$  is the length of the rightmost path of the tree pattern. We present efficient join operations on **occ** that help us to efficiently calculate the occurrence count of tree patterns from the occurrence count of their proper subtrees. Furthermore, we observed that in most of widely used real-world databases, while many vertices of a database tree have the same label, no two vertices on the same path are identically labeled. For this class of database trees, worst case space complexity of our algorithm is linear; a result comparable to the best existing results for per-tree frequency. We note that for such databases, worst case space complexity of the well-known existing algorithms for per-occurrence frequency, such as VTreeMiner [26] and MB3Miner [20], is still exponential. Based on the proposed subtree homeomorphism method, we develop an efficient pattern mining algorithm, called TPMiner. To evaluate the efficiency of TPMiner, we perform extensive experiments on both real-world and synthetic datasets. Our results show that TPMiner always outperforms most efficient existing algorithms such as VTreeMiner [26] and MB3Miner [20]. Furthermore, there are several cases where the improvement of TPMiner with respect to existing algorithm is very significant.

In Section 2, preliminaries and definitions related to the tree pattern mining problem are introduced. In Section 3, a brief overview on related work is given. In Section 4, we present the **occ** data-structure and our subtree homeomorphism algorithm. In Section 5, we introduce the TPMiner algorithm for finding frequent embedded tree patterns from rooted ordered trees. We empirically evaluate the effectiveness of TPMiner in Section 6. Finally, the paper is concluded in Section 7.




Fig. 1:  $T$  is a database tree where numbers next to vertices present preorder numbers;  $P1$  is a tree pattern where the vertices labeled by  $a$  and  $d$  form its rightmost path;  $P2$  is a rightmost path extension of  $P1$  where a new vertex labeled by  $a$  is added to an existing vertex of  $P1$  labeled by  $a$  as the rightmost child. Finally, dashed lines present a subtree homeomorphism mapping from  $P1$  to  $T$ .

## 2 Preliminaries

We assume the reader is familiar with the basic concepts in graph theory. The interested reader can refer to e.g., [12]. An *undirected (vertex-labeled) graph*  $G = (V, E, \lambda)$  consists of a vertex set  $V$ , an edge set  $E \subseteq \{e \subseteq V : |e| = 2\}$ , and a labeling function  $\lambda: V \rightarrow \Sigma$  which assigns a label from a finite set  $\Sigma$  to every vertex in  $V$ .  $G = (V, E, \lambda)$  is a *directed graph* if  $E$  is:  $E \subseteq V \times V$ , where  $\times$  is the *Cartesian product*. We use the notations  $V(G)$ ,  $E(G)$  and  $\lambda_G$  to refer to the set of vertices, the set of edges (or arcs) and the labeling function of  $G$ , respectively. The *size* of  $G$  is defined as the number of vertices of  $G$ . Two graphs  $G_1 = (V_1, E_1, \lambda_1)$  and  $G_2 = (V_2, E_2, \lambda_2)$  (either both are directed or both are undirected) are identical, written as  $G_1 = G_2$ , if  $V_1 = V_2$ ,  $E_1 = E_2$ , and  $\forall v \in V_1 : \lambda_1(v) = \lambda_2(v)$ . A *path* from a vertex  $v_0$  to a vertex  $v_n$  in a directed graph  $G = (V, E, \lambda)$  is a sequence of vertices such that  $\forall i, 0 \leq i \leq n-1, (v_i, v_{i+1}) \in E(G)$ . The *length* of a path is defined as its number of edges (number of vertices minus 1). A *cycle* is a path with  $v_0 = v_n$ .

An undirected graph not containing any cycles is called a *forest* and a connected forest is called a (*free*) *tree*. A *rooted tree* is a directed acyclic graph (DAG) in which: 1) there is a distinguished vertex, called *root*, that has no incoming edges, 2) every other vertex has exactly one incoming edge, and 3) there is an unique path from the root to any other vertex. In a rooted tree  $T$ ,  $u$  is the *parent* of  $v$  ( $v$  is the *child* of  $u$ ) if  $(u, v) \in E(T)$ . The transitive closure of the parent-child relation is called the *ancestor-descendant* relation. A *rooted ordered tree* is a rooted tree such that there is an order over the children of every vertex. Throughout this paper, we refer to *rooted ordered trees* simply as *trees*. *Preorder traversal* of a tree  $T$  is defined recursively as follows: first, visit  $root(T)$ ; and then for every child  $c$  of  $root(T)$  from left to right, perform a preorder traversal on the subtree rooted at  $c$ . The position of a vertex in the list of visited vertices during a preorder traversal is called its *preorder number*. We use  $p(v)$  to refer to the preorder number of vertex  $v$ . For instance, in Figure 1, numbers next to vertices of tree  $T$  present their preorder numbers. The *rightmost path*

of  $T$  is the path from  $root(T)$  to the last vertex of  $T$  visited in the preorder traversal. For example, in Figure 1, the vertices labeled by  $a$  and  $d$  form the rightmost path of  $P1$ . Two distinct vertices  $u$  and  $v$  are *relatives* if  $u$  is neither an ancestor nor a descendant of  $v$ . With  $p(u) < p(v)$ ,  $u$  is a left relative of  $v$ , otherwise, it is a right relative. For example, in tree  $T$  of Figure 1, vertices 1 and 4 are relatives and vertex 4 is a right relative of vertex 1.

A tree  $T$  is a *rightmost path extension* of a tree  $T'$  iff there exist vertices  $u$  and  $v$  such that: (i)  $\{v\} = V(T) \setminus V(T')$ , (ii)  $\{(u, v)\} = E(T) \setminus E(T')$ , (iii)  $u$  is on the rightmost path of  $T'$ , and (iv) in  $T$ ,  $v$  is a right relative of all children of  $u$ . We say that  $T$  is the rightmost path extension of  $T'$  with  $v$  attached at  $u$  and we denote  $T$  as  $RExtend(T', v, u)$ . For instance, in Figure 1,  $P2$  is a rightmost path extension of  $P1$ , where the new vertex  $v$  is labeled by  $a$ , the existing vertex  $u$  is also labeled by  $a$ , and  $v$  is added to  $u$  as the rightmost child.

A tree  $P$  is *subtree homeomorphic* to a tree  $T$  (denoted by  $P \preceq_h T$ ) iff there is a mapping  $\varphi : V(P) \rightarrow V(T)$  such that: (i)  $\forall v \in V(P) : \lambda_P(v) = \lambda_T(\varphi(v))$ , (ii)  $\forall u, v \in V(P) : u$  is the parent of  $v$  in  $P$  iff  $\varphi(u)$  is an ancestor of  $\varphi(v)$  in  $T$ , and (iii)  $\forall u, v \in V(P) : p(u) < p(v) \Leftrightarrow p(\varphi(u)) < p(\varphi(v))$ . For example, in Figure 1, dashed lines present a subtree homeomorphism mapping from  $P1$  to  $T$ . Under subtree isomorphism and homomorphism, the ancestor-descendant relationship between  $\varphi(u)$  and  $\varphi(v)$  in (ii) is strengthened into the parent-child relationship; under subtree homomorphism, (iii) is weakened into:  $p(u) < p(v) \Leftrightarrow p(\varphi(u)) \leq p(\varphi(v))$ , i.e., successive children of a vertex in the pattern can be mapped onto the same vertex in the database tree.  $P$  is *isomorphic* to  $T$  (denoted  $P \cong_i T$ ) iff  $P$  is subtree isomorphic to  $T$  and  $|V(P)| = |V(T)|$ .

Note that a pattern  $P$  under a subtree morphism can have several mappings to the same database tree  $T$ . When the matching operator is subtree homeomorphism, every mapping is called an *occurrence* (or *embedding*) of  $P$  in  $T$ . An occurrence (embedding) of a vertex  $v$  is an occurrence (embedding) of the pattern consisting of the single vertex  $v$ . The number of occurrences of  $P$  in  $T$  is denoted by  $NumOcc(P, T)$ .

Given a database  $\mathcal{D}$  consisting of trees and a tree  $P$ , *per-tree support* (or *per-tree frequency*) of  $P$  in  $\mathcal{D}$  is defined as:  $|\{T \in \mathcal{D} : P \preceq_h T\}|$ . *Per-occurrence support* (or *per-occurrence frequency*) of  $P$  in  $\mathcal{D}$  is defined as:  $\sum_{T \in \mathcal{D}} NumOcc(P, T)$ . In this paper, our focus is *per-occurrence support*. For the sake of simplicity, we use the term *support* (or *frequency*) instead of *per-occurrence support* (or *per-occurrence frequency*), and denote it by  $sup(P, \mathcal{D})$ .  $P$  is *frequent* ( $P$  is a *frequent embedded pattern*), iff its support is greater than or equal to a user defined integer threshold  $minsup > 0$ . The problem studied in this paper is as follows: given a database  $\mathcal{D}$  consisting of trees and an integer  $minsup$ , find every frequent pattern  $P$  such that  $sup(P, \mathcal{D}) \geq minsup$ .

We observe that when *per-occurrence support* is used, anti-monotonicity might be violated: it is possible that the support of  $P$  is greater than or equal to  $minsup$ , but it has a subtree whose support is less than  $minsup$ . For example, consider the database tree  $T$  of Figure 2 and suppose that  $minsup$  is 2. Then, while the pattern  $P1$  is infrequent, it has two frequent supertrees  $P2$  and  $P3$ . Therefore, in a more precise (and practical) definition, which is also used by algorithms such as VTreeMiner [26], tree  $P$  is frequent iff: 1)  $sup(P, \mathcal{D}) \geq minsup$ , and 2) the subtree  $P'$  generated by




Fig. 2: In the database tree  $T$  and for  $minsup = 2$ , while  $P1$  is infrequent, it has two frequent supertrees  $P2$  and  $P3$ .

removing the rightmost vertex of  $P$  is frequent. This means only frequent trees are extended to generate larger patterns.

### 3 Related work

Recently, many algorithms have been proposed for finding frequent embedded patterns from a database of tree-structured data that work with both *per-tree* and *per-occurrence* frequencies. Zaki presented VTreeMiner [26] to find embedded patterns from trees. For frequency counting he used an efficient data structure, called *scope-list*, and proposed rightmost path extension to generate non-redundant candidates. Later, he proposed the SLEUTH algorithm to mine embedded patterns from rooted unordered trees [25]. Tan et. al. [20] introduced the MB3Miner algorithm, where they use a unique occurrence list representation of the tree structure, that enables efficient implementation of their Tree Model Guided (TMG) candidate generation. TMG can enumerate all the valid candidates that fit in the structural aspects of the database. Chaoji et. al. [6] introduced a generic pattern mining approach that supports various types of patterns and frequency notions. A drawback of these algorithms is that in order to count the number of occurrences of a tree pattern  $P$  in a database tree  $T$ , they need to keep track of all occurrences of  $P$  in  $T$ . For example, in VTreeMiner, for every occurrence  $\varphi$  of  $P$  in  $T$  a separate element is stored in *scope-list*, that consists of the following components: (i) *TId* which is the identifier of the database tree that contains the occurrence, (ii)  $m$  which is  $\{\varphi(v) | v \in V(P) \setminus \{\text{rightmost vertex of } P\}\}$ , and (iii)  $s$  which is the scope of  $\varphi(u)$ , where  $u$  is the rightmost vertex of  $P$ . In the current paper, we propose a much more compact data-structure that can represent/encode all occurrences that have the rightmost path in common in  $O(d)$  space, where  $d$  is the length of the rightmost path of  $P$ . The number of such occurrences can be exponential. We then present efficient algorithms that calculate the occurrence count of tree patterns from the occurrence count of their proper subtrees.

There also exist several algorithms optimized for *per-tree frequency*. Zaki presented TreeMinerD [26] based on the *SV-list* data-structure. Tatikonda et. al. [22] developed TRIPS and TIDES using two sequential encodings of trees to systematically

generate and evaluate tree patterns. Wang et. al. [23] proposed XSpanner that uses a pattern growth-based method to find embedded tree patterns. The literature contains also many algorithms for mining patterns under subtree isomorphism; examples are in [1], [17], [2], [8] and [7]. From the applicability point of view, for instance in the prediction task, when longer range relationships are relevant, subtree homeomorphism becomes more useful than subtree isomorphism. For a more comprehensive study of related work, the interested reader can refer to e.g., [9].

A slightly different problem over rooted, ordered, and labeled trees is the tree inclusion problem: can a pattern  $P$  be obtained from a tree  $T$  by deleting vertices from  $T$ . In our terminology, is  $P$  present in  $T$  under subtree homeomorphism? Bille and Gørtz [3] recently proposed a novel algorithm that runs in linear space and sub-quadratic time, improving upon a series of polynomial time algorithms that started with the work in [15].

#### 4 Efficient tree mining under subtree homeomorphism

In this section, we present our method for subtree homeomorphism of rooted ordered trees. First in Section 4.1, we introduce the notion of *occurrence tree* and its rightmost path extension. Then in Section 4.2, we present the **occ-list** data-structure and, in Section 4.3, the operators for this data structure. In Section 4.4, we analyze space and time complexities of our proposed frequency counting method. We briefly compare our approach with other vertical frequency counting methods in Section 4.5.

##### 4.1 Occurrence trees and their extensions

Under rightmost path extension, a pattern  $P$  with  $k + 1$  vertices is generated from a pattern  $P'$  with  $k$  vertices by adding a vertex  $v$ , as the rightmost child, to a vertex in the rightmost path of  $P$ . Occurrences of  $P$  are rightmost path extensions of occurrences of  $P'$  with an occurrence of  $v$ . Therefore, an interesting way to construct occurrences of  $P$  is to look at the occurrences of  $v$  that can be a rightmost path extension of an occurrence of  $P'$ . First, we introduce the notion of *occurrence tree*. Then, we present the conditions under which a rightmost path extension of an occurrence tree yields another occurrence tree.

**Definition 1 (Occurrence tree)** Given an occurrence  $\varphi$  of  $P$  in  $T$ , we define the *occurrence tree*  $\mathcal{OT}(\varphi)$  as follows: (i)  $V(\mathcal{OT}(\varphi)) = \{\varphi(v) : v \in V(P)\}$ , (ii)  $root(\mathcal{OT}(\varphi)) = \varphi(root(P))$ , for every  $v \in V(\mathcal{OT}(\varphi))$ ,  $\lambda_{\mathcal{OT}(\varphi)}(v) = \lambda_T(v)$ , and (iii)  $E(\mathcal{OT}(\varphi)) = \{(\varphi(v_1), \varphi(v_2)) | (v_1, v_2) \in E(P)\}$ .

Notice, when  $(v_1, v_2) \in E(\mathcal{OT}(\varphi))$  and  $v_1$  is not the parent of  $v_2$  in  $T$ , then all intermediate vertices on the path from  $v_1$  to  $v_2$  are not part of  $V(\mathcal{OT}(\varphi))$ . For example, in Figure 3, the tree  $P$  has 3 occurrences in tree  $T$ .

Selecting a vertex not yet in the occurrence tree and performing a rightmost path extension does not always result in another occurrence tree. Proposition 1 below lists properties that hold when the rightmost path extension is an occurrence tree.




Fig. 3: From left to right, a database tree  $T$ , a pattern  $P$  and three occurrence trees  $OT(\varphi_1)$ ,  $OT(\varphi_2)$ , and  $OT(\varphi_3)$ . Labels are inside vertices, preorder numbers are next to vertices. The occurrence trees are represented by showing the occurrences of the pattern vertices in bold. Their edges are the images of the edges in the pattern; for example,  $OT(\varphi_3)$  refers to the tree formed by vertices 0, 4 and 5, with 0 as the root, and with  $(0, 4)$  and  $(0, 5)$  as edges.




Fig. 4: A database tree is shown 4(a), together with four rightmost path extensions of different occurrence trees in that database tree. However, only 4(b) is itself an occurrence tree in the database tree; 4(c), 4(d) and 4(e) violate conditions (ii), (iii) and (i) of Proposition 2, respectively.

**Proposition 1** *Let  $\varphi'$  be an occurrence of a pattern  $P'$  in a database tree  $T$  and  $OT' = OT(\varphi')$ . Let  $x$  be a vertex of  $T$  outside  $OT'$ , and  $y$  a vertex on the rightmost path of  $OT'$ . If  $OT = RExtend(OT', x, y)$  is an occurrence tree in  $T$ , then: (i)  $root(OT')$  is an ancestor of  $x$  in  $T$ , (ii) of all ancestors of  $x$  in  $T$  that belong to  $OT'$ ,  $y$  is the largest one in the preorder over  $T$ , and (iii) for each vertex  $w$  in  $OT'$  such that  $p(w) > p(y)$ ,  $w$  is a left relative of  $x$  in  $OT$  and in  $T$ .*

*Proof* By the definition of rightmost path extension presented in Section 2, either  $y$  is  $root(OT')$  or it is a descendant of  $root(OT')$ ; moreover  $y$  is the parent of  $x$ , hence (i) holds. As  $y$  is on the rightmost path in  $OT'$  and  $y$  is the parent of  $x$  in  $OT$  (and children of  $y$  are relatives of  $x$  in  $T$ ),  $y$  is, among the ancestors of  $x$  in  $T$ , the largest one that belongs to  $OT'$  (ii). If  $p(w) > p(y)$  for a vertex  $w \in V(OT')$ , since  $y$  is on the rightmost path of  $OT'$ ,  $w$  is a descendant of  $y$ . Furthermore, since  $x$  is the rightmost child of  $y$  in  $OT$ ,  $w$  is a left relative of  $x$  in  $OT$  and in  $T$  (iii). ■

As an example of Proposition 1, consider Figure 4, where 4(a) presents a database tree and 4(b) shows that an occurrence tree  $OT'$  consisting of vertices 0 and 1 is extended to another occurrence tree  $OT$  consisting of vertices 0, 1 and 5. Vertex 0 is an ancestor of vertex 5 in the database tree (condition (i)); among all ancestors of vertex 5 in the database tree that belong to  $OT'$ , vertex 0 is the largest one in the



preorder over the database tree (condition (ii)); and vertex 1 is a left relative of vertex 5 in  $OT$  and in the database tree (condition (iii)).

Let  $P'$  be a tree pattern. The next proposition lists the conditions that are sufficient to ensure that a rightmost path extension of an occurrence tree  $OT(\varphi')$  of  $P'$  with an edge  $(y, x)$  yields an occurrence tree of another tree pattern  $P$ , where  $P$  is a rightmost path extension of  $P'$ .

**Proposition 2** *Let  $\varphi'$  be an occurrence of a pattern  $P'$  in a database tree  $T$  and  $OT' = OT(\varphi')$ . Let  $x$  be a vertex of  $T$  outside  $OT'$ , and  $y$  a vertex on the rightmost path of  $OT'$ . If: (i)  $y$  is an ancestor of  $x$  in the database tree  $T$ , (ii) of all vertices of  $OT'$  that are ancestors of  $x$  in the database tree  $T$ ,  $y$  is the largest one, and (iii)  $p(x) > p(w)$  for all  $w \in OT'$ , then  $RExtend(OT', x, y)$  is an occurrence tree. For a given vertex  $x$  and an occurrence tree  $OT$ , if there exists a vertex  $y$  such that  $RExtend(OT, x, y)$ ,  $y$  is unique.*

*Proof* Let  $u$  be the vertex of  $P'$  such that  $y = \varphi'(u)$ . To define  $P$ , we set  $V(P) = V(P') \cup \{v\}$  with  $v$  a new vertex with the same label as  $x$ , and  $E(P) = E(P') \cup \{(u, v)\}$  such that  $v$  is the rightmost child of  $u$ . Define  $OT$  as  $V(OT) = V(OT') \cup \{x\}$  and  $E(OT) = E(OT') \cup \{(y, x)\}$  and set  $\varphi(P) = \varphi'(P') \cup \{v \rightarrow x\}$ . By the construction and the assumptions,  $OT = RExtend(OT', x, y)$ . We show  $OT$  is a occurrence tree of  $P$  in  $T$ . From (i) and (ii) it follows that  $x$  is on the rightmost path from  $root(OT)$  and it is a descendant of  $y$  in  $T$  and from (iii) that the children of  $y$  in  $OT$  are left relatives of  $x$ , hence,  $\varphi$  is an embedding of  $P$  in  $T$  and  $OT$  is an occurrence tree of  $P$  in  $T$ . We note since all vertices of a tree have a unique preorder number,  $y$  is unique, if it exists. ■

Figure 4 shows a database tree (4(a)) and four rightmost path extensions of occurrence trees; however, only one of these extensions is another occurrence tree (4(b)), the other ones violate the conditions of Proposition 2.

To turn the conditions of Proposition 2 into a practical method, we need a compact way to store occurrence trees and an efficient way to check the conditions. For the latter, we take advantage of the solution of Dietz [13]. He has designed a numbering scheme based on tree traversal order to determine the ancestor-descendant relationship between any pair of vertices. This scheme associates each vertex  $v$  with a pair of numbers  $\langle p(v), p(v) + size(v) \rangle$ , where  $size(v)$  is an integer with certain properties (which are met when e.g.,  $size(v)$  is the number of descendants of  $v$ ). Then, for two vertices  $u$  and  $v$  in a given database tree,  $u$  is an ancestor of  $v$  iff  $p(u) < p(v)$  and  $p(v) + size(v) \leq p(u) + size(u)$ , and  $v$  is a right relative of  $u$  iff  $p(u) + size(u) < p(v)$ . In several algorithms, such as VTreeMiner and TreeMinerD [26] and TwigList [18], variants of this scheme have been used to design efficient methods for counting occurrences of tree patterns based on occurrences of their subtrees.

Our contribution is to introduce data structures, based on the Dietz numbering scheme, that allow us to speed up counting of all occurrences of tree patterns. For example, while the algorithm of [26] keep track of all occurrences, we only store the occurrences that have distinct rightmost paths. We start with introducing some additional notations. The scope of a vertex  $x$  in a database tree  $T$ , denoted  $x.scope$ ,

is a pair  $(l, u)$ , where  $l$  is the preorder number of  $x$  in  $T$  and  $u$  is the preorder number of the rightmost descendant of  $x$  in  $T$ . We use the notations  $x.scope.l$  and  $x.scope.u$  to refer to  $l$  and  $u$  of the scope of  $x$ .

**Definition 2 (rdepth)** Let  $x$  be a vertex on the rightmost path of a tree  $T$ . The *rdepth* of  $x$  in  $T$ , denoted  $rdep_T(x)$ , is the length of the path from the root of  $T$  to  $x$ .

A vertex  $x$  on the rightmost path of  $T$  is uniquely distinguished by  $rdep_T(x)$ . For example in Figure 4(a), the *rdepth* of vertices 0, 3 and 5 is 0, 1 and 2, respectively (and for the other vertices, *rdepth* is undefined).

**Proposition 3** Let  $OT'$  be an occurrence tree of a tree pattern  $P'$  in a database tree  $T$ ,  $x \in V(T) \setminus V(OT')$  and  $y$  a vertex on the rightmost path of  $OT'$  but not the rightmost vertex (i.e., the rightmost path of  $OT'$  has a vertex  $z$  such that  $rdep_{OT'}(z) = rdep_{OT'}(y) + 1$ ). We have:  $RExtend(OT', x, y)$  is an occurrence tree iff

$$z.scope.u < x.scope.l \leq y.scope.u \quad (1)$$

*Proof* First, assume  $RExtend(OT', x, y)$  is an occurrence tree. By Proposition 1,  $y$  is the largest vertex of  $OT'$  that is an ancestor of  $x$ , hence, in the database tree  $T$ , the tree rooted at  $x$  is a subtree of the tree rooted at  $y$ . That means

$$y.scope.l < x.scope.l \leq x.scope.u \leq y.scope.u \quad (2)$$

Vertex  $z$  and all vertices in the subtree of  $z$  are left relatives of  $x$  and have a preorder number smaller than that of  $x$ . Hence  $z.scope.u < x.scope.l$ . Combining with Inequation 2, we obtain  $z.scope.u < x.scope.l \leq y.scope.u$ .

For the other direction, Inequation 1 yields that  $x$  is a right relative of  $z$  and it is in the scope of the subtree of  $y$ , hence,  $y$  is an ancestor of  $x$  ((i) of Proposition 2) and  $z$  is not an ancestor of  $x$ , so  $y$  is the largest ancestor of  $x$  that belongs to  $OT'$  ((ii) of Proposition 2). Also, the preorder number of  $x$  is larger than that of any vertex in the subtree of  $z$  and hence of any vertex in  $OT'$  ((iii) of Proposition 2). Hence, by Proposition 2,  $RExtend(OT', x, y)$  is an occurrence tree. ■

**Proposition 4** Let  $OT'$  be an occurrence tree of a tree pattern  $P'$  in a database tree  $T$ ,  $x \in V(T) \setminus V(OT')$  and  $y$  the rightmost vertex of  $OT'$ . We have:  $RExtend(OT', x, y)$  is an occurrence tree iff

$$y.scope.l < x.scope.l \text{ and } x.scope.u \leq y.scope.u \quad (3)$$

*Proof* First, assume  $RExtend(OT', x, y)$  is an occurrence tree. Similar to the proof of Proposition 3, by Proposition 1, we get

$$y.scope.l < x.scope.l \leq x.scope.u \leq y.scope.u \quad (4)$$

For the other direction, we assume  $y.scope.l < x.scope.l$  and  $x.scope.u \leq y.scope.u$ . This implies that the tree rooted at  $x$  is a subtree of the tree rooted at  $y$  and that the preorder number of  $x$  is larger than the preorder number of  $y$  and hence that all conditions of Proposition 2 are satisfied, and  $RExtend(OT', x, y)$  is an occurrence tree. ■

For example, in Figure 4, first let  $OT'$  refer to the occurrence tree consisting of a single vertex 0. The scopes of vertices 0 and 1 are  $(0, 5)$  and  $(1, 1)$ , respectively. The lower bound of the scope of vertex 1 is greater than the the lower bound of the scope of vertex 0; and the upper bound of the scope of vertex 1 is smaller than or equal to the upper bound of the scope of vertex 0. Therefore, Inequation 3 holds and  $RExtend(OT', 1, 0)$  is an occurrence tree. Now, let  $OT'$  refer to the occurrence tree consisting of vertices 0 and 1. The scope of vertex 5 is  $(5, 5)$ . The lower bound of the scope of vertex 5 is greater than the upper bound of the scope of vertex 1; and it is smaller than or equal to the upper bound of the scope of vertex 0. Hence, Inequation 1 holds and  $RExtend(OT', 5, 0)$  is an occurrence tree.

#### 4.2 **Occ-list**: an efficient data structure for tree mining under subtree homeomorphism

For the rightmost path extension of an occurrence tree, it suffices to know its rightmost path. Different occurrences of a pattern can have the same rightmost path. The key improvement over previous work is that we only store information about the rightmost path of occurrence trees and that different occurrence trees with the same rightmost path are represented by the same data element. All occurrences of a pattern in a database tree are represented by **occ-list**, a list of occurrences. An element **occ** of **occ-list** represents all occurrences with a particular rightmost path. The element has four components:

- *TId*: the identifier of the database tree that contains the occurrences represented by **occ**.
- *scope*: the scope in the database tree *TId* of the last vertex in the rightmost path of the occurrences represented by **occ**,
- *RP*: an array containing the upper bounds of the scopes of the vertices in the rightmost path of the occurrences represented by **occ**, i.e., with  $x$  the vertex at rdepth  $j$  in the rightmost path of the pattern  $P$  and  $\varphi$  one of the occurrences represented by **occ**,  $RP[j] = \varphi(x).scope.u$ ; note that this is the same value for all occurrences  $\varphi$  represented by **occ**<sup>1</sup>.
- *multiplicity*: the number of occurrences represented by **occ**.

For all occurrences of a pattern  $P$  that have the same *TId*, *scope* and *RP*, one **occ** in the **occ-list** of  $P$  is generated and its *multiplicity* shows the number of such occurrences. All **occs** of a pattern of size 1 have *multiplicity* 1 as their (single vertex) rightmost paths are all different. Every occurrence is represented by exactly one **occ**. We refer to the **occ-list** of  $P$  by **occ-list**( $P$ ). It is easy to see that the frequency of  $P$  is equal to  $\sum_{oc \in \text{occ-list}(P)} oc.multiplicity$ . An example of **occ-list** is shown in Figure 5.

<sup>1</sup> The upper bound of the scope of the last vertex is already available in *scope*; for convenience of presentation, the information is duplicated in *RP*.




Fig. 5: An example of **occ-list**.  $T_0$  and  $T_1$  are two database trees and minimum-support is equal to 2. The figure presents the **occ-lists** of some frequent 1-tree patterns, frequent 2-tree patterns, frequent 3-tree patterns and frequent 4-tree patterns.

#### 4.3 Operations on the **occ-list** data structure

Let  $P$  be a tree of size  $k + 1$  generated by adding a vertex  $v$  to a tree  $P'$  of size  $k$ . There are two cases:

1.  $v$  is added to the rightmost vertex of  $P'$ . We refer to this case as *leaf join*.
2.  $v$  is added to a vertex in the rightmost path of  $P'$ , but not its rightmost vertex. We refer to this case as *inner join*.

They correspond to Propositions 3 and 4.

**Proposition 5 (leaf join)** Let  $v$  be a one element pattern,  $P'$  a pattern with  $u$  as the rightmost vertex,  $d = rdep_{P'}(u)$  and  $P = RExtend(P', v, u)$ . Let **occ-list**( $P'$ ), **occ-list**( $v$ ) and **occ-list**( $P$ ) be the representations of the occurrences of  $P'$ ,  $v$  and  $P$ , respectively. We have:  $oc \in \text{occ-list}(P)$  iff there exists an  $oc' \in \text{occ-list}(P')$  and an  $ov \in \text{occ-list}(v)$  such that

- (i)  $oc'.TId = ov.TId = oc.TId$  (all occurrences are from the same database tree),
- (ii)  $oc'.scope.l < ov.scope.l$  and  $ov.scope.u \leq oc'.scope.u$ ,
- (iii)  $oc.RP[i] = oc'.RP[i]$  ( $0 \leq i \leq d$ ) and  $oc.RP[d+1] = ov.scope.u$  (i.e., copy of  $oc'.RP$  and an extra element),
- (iv)  $oc.scope = ov.scope$ , and
- (v)  $oc.multiplicity = oc'.multiplicity$

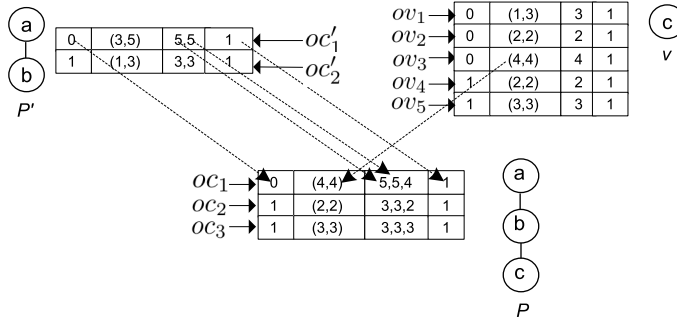


Fig. 6: Details of the relationship between  $\text{occ-list}(P')$ ,  $\text{occ-list}(v)$  and  $\text{occ-list}(P)$  for the database trees of Figure 5. The entries  $oc'_1$ ,  $ov_3$  and  $oc_1$  satisfy the properties of Proposition 5. Also the tuples  $(oc'_2, ov_4, oc_2)$  and  $(oc'_2, ov_5, oc_3)$  satisfy the properties. The proposition is exploited in Algorithm 1 below. Its *leaf-join* operation uses  $\text{occ-list}(P')$  and  $\text{occ-list}(v)$  to compute  $\text{occ-list}(P)$ .

*Proof* First, assume  $oc \in \text{occ-list}(P)$ , hence it represents  $oc.\text{multiplicity}$  occurrence trees of pattern  $P$  in database tree  $oc.TId$ . Each of these occurrence trees share the same rightmost path. Hence they can be decomposed into occurrence trees of pattern  $P'$  sharing the same rightmost path and a particular occurrence of  $v$ . The latter is represented by an element  $ov$  of  $\text{occ-list}(v)$ . The formers are represented by an element  $oc'$  of  $\text{occ-list}(P')$ . Now we show that conditions (i), ..., (v) hold between  $oc$  and these elements  $oc'$  and  $ov$ . Condition (i) holds because all occurrences are in database tree  $oc.TId$ , (ii) follows from Proposition 4, condition (iii) follows because the rightmost path of the occurrence trees represented by  $oc'$  is identical to that of  $oc$ , except for the last element which is removed, (iv) follows from the definition of *scope*, and (v) holds because  $oc$  and  $oc'$  represent the same number of occurrences.

Second, assume there are an  $oc' \in \text{occ-list}(P')$  and an  $ov \in \text{occ-list}(v)$  such that conditions (i), ..., (v) hold. From (i) it follows that  $oc'$  and  $ov$  present occurrences in the same database tree. Because (ii) holds, it follows from Proposition 4 that all occurrence trees of  $P'$  represented by  $oc'$  can be extended with  $ov$  into occurrence trees of  $P$ ; all these occurrence trees have the same rightmost path and have  $ov$  as their rightmost vertex, moreover they are the only ones in the database tree  $oc'.TId$  with such a rightmost path. Hence, the element  $oc$  that satisfies properties (i), ..., (v) is indeed an element of  $\text{occ-list}(P)$  as has to be proven. ■

Figure 6 illustrates the proposition. The proposition is the basis for the *leaf-join* operation in Algorithm 1 below.

In contrast with leaf join, which is performed between one **occ** of a tree pattern and one **occ** of a vertex, inner join is performed between a set of **occs** of a tree pattern and one **occ** of a vertex.

**Proposition 6 (Inner join)** Let  $v$  be a one element pattern,  $P'$  a pattern,  $u$  a vertex on the rightmost path of  $P'$  but not the rightmost vertex,  $c = \text{rdep}_{P'}(u)$  and

$P = R\text{Extend}(P', v, u)$ . Let  $\mathbf{occ-list}(P')$ ,  $\mathbf{occ-list}(v)$  and  $\mathbf{occ-list}(P)$  be the representations of the occurrences of  $P'$ ,  $v$  and  $P$ , respectively. We have:  $oc \in \mathbf{occ-list}(P)$  iff there exists a subset  $oc'_1, \dots, oc'_m$  ( $m \geq 1$ ) of  $\mathbf{occ-list}(P')$  and an  $ov \in \mathbf{occ-list}(v)$  such that

- (i)  $ov.TId = oc'_1.TId = \dots = oc'_m.TId = oc.TId$  (all occurrences are from the same database tree),
- (ii)  $oc'_i.RP[k] = oc'_j.RP[k]$ , for all  $i, j \in [1..m]$  and for all  $k \in [0..c]$ ,
- (iii)  $oc'_i.RP[c+1] < ov.scope.l \leq oc'_i.RP[c]$  for all  $i \in [1..m]$ ,
- (iv)  $oc'_1, \dots, oc'_m$  is maximal with the conditions (i)-(iii),
- (v)  $oc.RP[i] = oc'.RP[i]$ ,  $0 \leq i \leq c$ , and  $oc.RP[c+1] = ov.scope.u$  (copy of part of  $oc'.RP$  and an extra element  $ov.scope.u$ ),
- (vi)  $oc.scope = ov.scope$ , and
- (vii)  $oc.multiplicity = \sum_{i=1}^m oc'_i.multiplicity$ .

*Proof* First, assume  $oc \in \mathbf{occ-list}(P)$ , hence it represents  $oc.multiplicity$  occurrence trees of pattern  $P$  in database tree  $oc.TId$ . All these occurrence trees share the same rightmost path. Hence, they can be decomposed into occurrence trees of pattern  $P'$  sharing the first  $c+1$  vertices of the rightmost path and a particular occurrence of  $v$ . The latter is represented by an element  $ov$  of  $\mathbf{occ-list}(v)$ . The formers are represented by elements  $oc'_1, \dots, oc'_m$  of  $\mathbf{occ-list}(P')$ . Now we show that conditions (i), ..., (vii) hold between  $oc$  and these elements  $oc'_1, \dots, oc'_m$  and  $ov$ . Condition (i) holds because all occurrence trees are in database tree  $oc.TId$ , (ii) holds because all occurrence trees represented by  $oc'_1, \dots, oc'_m$  share the first  $c+1$  vertices of the rightmost paths, (iii) and (iv) follow from Proposition 3, (v) follows because the first  $c+1$  vertices of the rightmost paths of the occurrence trees represented by  $oc$  are identical to those of  $oc'_1, \dots, oc'_m$ , and the rightmost vertex of the occurrence trees represented by  $oc$  is the vertex represented by  $ov$ , (vi) follows from the definition of  $scope$ , and (vii) holds because  $oc$  represents  $\sum_{i=1}^m oc'_i.multiplicity$  occurrences.

Second, assume there are a maximal subset  $oc'_1, \dots, oc'_m$  of  $\mathbf{occ-list}(P')$  and an  $ov \in \mathbf{occ-list}(v)$  such that conditions (i), ..., (vii) hold. From (i) it follows that  $oc'$  and  $ov$  present occurrences in the same database tree. Because (iii) and (iv) hold, it follows from Proposition 3 that all occurrence trees of  $P'$  represented by  $oc'_1, \dots, oc'_m$  can be extended with  $ov$  into occurrence trees of  $P$ ; all these extended occurrence trees have the same rightmost path and have  $ov$  as their rightmost vertex, moreover they are the only ones in the database tree  $oc'.TId$  with such a rightmost path. Hence, the element  $oc$  that satisfies properties (i), ..., (vii) is indeed an element of  $\mathbf{occ-list}(P)$  as has to be proven. ■

Figure 7 illustrates the proposition. The proposition is the basis for the *inner\_join* operation in Algorithm 1 below.

#### 4.4 Complexity analysis

*Space complexity.* Given a database  $\mathcal{D}$ , space complexity of the  $\mathbf{occ-list}$  of a pattern of size 1 is  $O(n \times |\mathcal{D}|)$ , where  $n$  is the maximum number of vertices that a database




Fig. 7: Details of the relationship between  $\text{occ-list}(P')$ ,  $\text{occ-list}(v)$  and  $\text{occ-list}(P)$  for the database trees of Figure 5. The entries  $\{oc'_1, oc'_2\}$ ,  $ov_1$  and  $oc$  satisfy the properties of Proposition 6. As  $c = 0$ , rightmost paths of the occurrence trees represented by  $oc'_1$  and  $oc'_2$  share the first vertex, that is vertex 0 of  $T_0$ ; and rightmost paths of the occurrence trees represented by  $oc$  share the first and second vertices, that are vertices 0 and 3 of  $T_0$ . The proposition is exploited in Algorithm 1 below. Its *inner\_join* operation uses  $\text{occ-list}(P')$  and  $\text{occ-list}(v)$  to compute  $\text{occ-list}(P)$ .




Fig. 8: Figure (a) shows a tree pattern  $P$  and (b) a database tree  $T$ . The number of occurrences of  $P$  in  $T$  is exponential in terms of  $n$  and  $k$ . In this case, the size of  $\text{occ-list}$  is linear, however, the size of the data-structures generated by the other algorithms is exponential.

tree  $T \in \mathcal{D}$  has. For larger patterns, space complexity of the  $\text{occ-list}$  of a pattern  $P$  is  $O(b \times d \times |\mathcal{D}|)$ , where  $b$  is the maximum number of occurrences with distinct rightmost paths that  $P$  has in a database tree  $T \in \mathcal{D}$ , and  $d$  is the length of the rightmost path of  $P$ .

Compared to data-structures generated by other algorithms such as VTreeMiner,  $\text{occ-list}$  is often substantially more compact. Given a database  $\mathcal{D}$ , the size of the data-structure generated by VTreeMiner for a pattern  $P$  is  $O(e \times k \times |\mathcal{D}|)$ , where  $e$  is the maximum number of occurrences that  $P$  has in a database tree  $T \in \mathcal{D}$  and  $k$  is  $|V(P)|$ . We note that on one hand,  $k \geq d$  and the other hand,  $e \geq b$ . In particular,  $e$  can be significantly larger than  $b$ , e.g., it can be exponentially (in terms of  $n$  and




Fig. 9: Figure (a) shows a tree pattern  $P$  and (b) a database tree  $T$ . The number of occurrences of  $P$  in  $T$  is exponential in terms of  $n$  and  $k$ . In this case, the size of the **occ-list** and the size of the data-structures generated by the other algorithms for  $P$  are exponential.

$k$ ) larger than  $b$ . Therefore, in the worst case, the size of the data-structure generated by VTreeMiner is exponentially larger than **occ-list** (and it is never smaller than **occ-list**). An example of this situation is shown in Figure 8. In this figure, pattern  $P$  has  $\binom{n-1}{k-1} \geq \left(\frac{n-1}{k-1}\right)^{k-1} \geq \left(\frac{n}{k}\right)^{k/2}$  occurrences in the database tree  $T$  ( $k > 2$  and  $n \geq k$ ). If  $k = (n+1)/2$ , the size of the data-structure generated by VTreeMiner will be  $\Omega(2^{n/4})$ . However, in this case, the size of the **occ-list** generated for  $P$  is  $\Omega(n)$ . More precisely, **occ-list**( $P$ ) has  $n - k + 1$  **occs**, where for each  $i$ ,  $k - 1 \leq i \leq n - 1$ , there exists an **occ** with  $TId = 0$ ,  $scope = (i, i)$ ,  $RP = \{n - 1, i\}$  and  $multiplicity = \binom{i-1}{k-2}$ .

We note that there are cases where the size of **occ-list** (as well as the size of the data-structures used by the other algorithms) becomes exponential. An example of this situation is presented in Figure 9, where for  $k = n/2$ , the size of **occ-list** as well as the size of *scope-list* used by VTreeMiner become exponential. More precisely, for each  $k$ -element combination of the vertices of the database tree, there exists an **occ** in the **occ-list** of the tree pattern, where  $TId$  is 0, the lower bound of *scope* is the preorder number of the vertex with the largest depth in the combination, the upper bound of *scope* is  $n - 1$ ,  $RP$  is an array of size  $k$  filled by  $n - 1$ , and *multiplicity* is 1.

In most of real-world databases, such as CSLOGS [26] and NASA [5], while several vertices of a database tree have the same label, no two vertices on the same path are identically labeled. For trees with this property, while worst case space complexity of **occ-list** is linear, worst case size of *scope-list* remains exponential.

*Time complexity.* We study time complexity of leaf join and inner join:

- A leaf join between two **occs** takes  $O(d)$  time with  $d$  the length of the rightmost path of  $P$ . Since a pattern larger than 1 has  $O(b \times |\mathcal{D}|)$  **occs** and a pattern of size 1 has  $O(n \times |\mathcal{D}|)$  **occs** and leaf join is performed between every pairs of **occs** with the same  $TId$ , worst case time complexity of leaf join between two **occ-lists** will be  $O(d \times b \times n \times |\mathcal{D}|)$ .



- In the inner join of the **occ-lists** of a tree pattern  $P'$  and a vertex  $v$ , given an **occ**  $ov$  of  $v$ , it takes  $O(h \times d)$  time to find subsets of the **occ-list** of  $P'$  that satisfy the conditions of Proposition 6, where  $h$  is the number of **occs** in the **occ-list** of  $P'$ . We note that **occs** of an **occ-list** can be automatically sorted first based on their  $TIds$  and second, based on their  $RPs$ . This makes it possible to find the subsets of the **occ-list** of  $P'$  that satisfy the conditions of Proposition 6 only by one scan of the **occ-list** of  $P'$ . During this scan, it is checked whether: (i) the current and the previous **occs** have the same  $TId$ , (ii) the current and the previous **occs** have the same  $RP[0], \dots, RP[c]$ , (iii)  $RP[c+1]$  of the current **occ** is less than  $ov.scope.l$ , and (iv)  $RP[c]$  of the current **occ** is greater than or equal to  $ov.scope.l$ . During the scan of **occ-list**( $P'$ ), after finding a subset  $S$  that satisfies the conditions of Proposition 6, it takes  $O(d)$  time to perform the inner join of  $S$  and  $ov$ . As a result, it takes  $O(h \times d)$  time to perform the inner join of the **occ-list** of  $P'$  and an **occ** of  $v$ . Since  $h$  is  $O(b \times |\mathcal{D}|)$  and  $v$  has  $O(n \times |\mathcal{D}|)$  **occs**, and since inner join is done between every pairs of subsets  $S$  and **occs**  $ov$  that have the same  $TIds$ , time complexity of inner join will be  $O(d \times b \times n \times |\mathcal{D}|)$ .

Therefore, frequency of a pattern can be counted in  $O(d \times b \times n \times |\mathcal{D}|)$  time. We note that in VTreeMiner, frequency of a pattern is counted in  $O(k \times e^2 \times |\mathcal{D}|)$  time (recall  $k \geq d$  and  $e \geq b$ , in particular, it is possible that  $e \gg b$ ).

#### 4.5 A brief comparison with other vertical frequency counting approaches

As mentioned before, algorithms such as VTreeMiner [26] and MB3Miner [20] need to keep track of all occurrences. Obviously, our approach is more efficient as it simultaneously represents and processes all occurrences sharing the rightmost path in  $O(d)$  space, where  $d$  is the length of the rightmost path of the pattern. TreeMinerD [26], developed by Zaki for computing *per-tree frequency*, is more similar to our approach as it also processes rightmost paths of occurrences. However, there are significant differences. First, while TreeMinerD computes only *per-tree frequency*, our algorithm performs a much more expensive task and computes *per-occurrence frequency*. Second, TreeMinerD applies different join strategies.

### 5 TPMiner: an efficient algorithm for finding frequent embedded tree patterns

In this section, we first introduce the TPMiner algorithm and then, we discuss an optimization technique used to reduce the number of generated infrequent patterns.

#### 5.1 The TPMiner algorithm

Having defined the operations of leaf join and inner join and having analyzed their properties, we can now introduce our tree pattern miner, called TPMiner (**T**ree **P**attern **M**iner). TPMiner builds all frequent patterns and maintains the rightmost paths of all their occurrences.

**Algorithm 1** High level pseudo code of the TPMiner algorithm.

---

```

1: TPMiner
2: Input:  $\mathcal{D}$  {a set of database trees},  $minsup$  {the minimum support threshold}
3: Output:  $\mathcal{P}$  {the set of frequent patterns}
4: Compute the set  $\mathcal{P}_1$  of frequent patterns of size 1 along with their occ-lists
5:  $\mathcal{P} \leftarrow \mathcal{P}_1$ 
6: for each  $P$  in  $\mathcal{P}_1$  do
7:    $Extend(P, \mathcal{P}_1, minsup, \mathcal{P})$ 
8: end for

```

---

```

1:  $Extend(P, \mathcal{P}_1, minsup, \mathcal{P})$ 
2: Input:  $P$  {a frequent pattern},  $\mathcal{P}_1$  {the set of frequent patterns of size 1},  $minsup$  {the minimum support threshold}
3: Input and Output:  $\mathcal{P}$  {the set of frequent patterns}
4: Side effect:  $\mathcal{P}$  is updated with frequent rightmost path extensions of  $P$ 
5: for each  $P_1$  in  $\mathcal{P}_1$  do
6:   {Let  $d$  be the length of the rightmost path of  $P$ }
7:   for  $i = 0$  to  $d$  do
8:     {Let  $u$  be the vertex of  $P$  such that  $rdep_P(u) = i$ }
9:      $P_n \leftarrow RExtend(P, P_1, u)$ 
10:    if  $i = d$  then
11:       $occ\_list(P_n) \leftarrow leaf\_join(P, P_1)$ 
12:    else
13:       $occ\_list(P_n) \leftarrow inner\_join(P, P_1)$ 
14:    end if
15:    if  $sup(P_n, \mathcal{D}) \geq minsup$  then
16:      Add  $P_n$  to  $\mathcal{P}$ 
17:       $Extend(P_n, \mathcal{P}_1, minsup, \mathcal{P})$ 
18:    end if
19:  end for
20: end for

```

---

TPMiner follows a depth-first strategy for generating tree patterns. First, it extracts frequent patterns of size 1 (frequent vertex labels) and constructs their **occ-lists**. This step can be done by one scan of the database. Every **occ** of a pattern of size 1 represents one occurrence of the pattern, where its *RP* contains the upper bound of the scope of the occurrence, its *scope* contains the scope of the occurrence, and its *multiplicity* is 1. Then, larger patterns are generated using rightmost path extension. For every tree  $P$  of size  $k + 1$  ( $k \geq 1$ ) which is generated by adding a vertex  $v$  to a vertex on the rightmost path of a tree  $P'$ , the algorithm computes the **occ-list** of  $P$  by joining the **occ-lists** of  $P'$  and  $v$ . If  $v$  is added to the the rightmost vertex of  $P'$ , a *leaf join* is performed; otherwise, an *inner join* is done. The high level pseudo code of TPMiner is given in Algorithm 1.  $\mathcal{P}$  is used to store all frequent patterns. Every tree pattern is generated in  $O(1)$  time, hence, time complexity of Algorithm 1 is  $O(d \times b \times n \times |\mathcal{D}| \times \mathcal{C})$ , where  $\mathcal{C}$  is the number of generated candidates.

## 5.2 An optimization technique for candidate generation

In rightmost path extension, a new tree  $P$  is generated by adding a new vertex to some vertex on the rightmost path of an existing frequent pattern  $P'$ , therefore, it is already

Table 1: Summary of real-world datasets.

Dataset	# Transactions	# Vertices	# Vertex labels	Transaction size	
				Maximum	Average
CSLOGS32241	32, 241	240, 716	10, 698	435	13.9323
Prions	17, 551	227, 203	111	37	24.9497
NASA	1, 000	163, 753	333	471	326.506

known that  $P$  has (at least) one frequent subtree. In an improved method proposed in [26], to generate the tree  $P$ , two patterns  $P1$  and  $P2$  such that their subtrees induced by all but the rightmost vertices are the same, are merged. In this merge, the rightmost vertex of  $P2$  (which is not in  $P1$ ) is added to  $P1$  as the new vertex. In this way, we already know that  $P$  has (at least) two subtrees that are frequent, therefore, trees that have only one frequent subtree are not generated. This can reduce the number of trees that are generated but are infrequent.

## 6 Experimental Results

We performed extensive experiments to evaluate the efficiency of the proposed algorithm, using data from real applications as well as synthetic datasets. The experiments were done on one core of a single AMD Processor 270 clocked at 2.0 GHz with 16 GB main memory and  $2 \times 1$  MB L2 cache, running Ubuntu Linux 12.0. The program was compiled by the GNU C++ compiler 4.0.2.

VTreeMiner (also called TreeMiner) [26] is a well-known algorithm for finding all frequent embedded patterns from trees. Therefore, we select this algorithm for our comparisons. Recently more efficient algorithms, such as TreeMinerD [26], XSpanner [23], TRIPS and TIDES [22], have been proposed for finding frequent embedded tree patterns. However, they only compute the per-tree frequency instead of the per-occurrence frequency. Some other algorithms, such as [24] and [16], find maximal embedded patterns which are a small subset of all frequent tree patterns. To the best of our knowledge at the time of writing this paper, MB3Miner [20] is the most efficient recent algorithm for finding frequent embedded tree patterns. MB3Miner works with both per-tree frequency and per-occurrence frequency. Here, we use the version that works with the per-occurrence frequency. We note MB3Miner generates only the frequent patterns such that all subtrees are frequent, therefore, it might produce fewer frequent patterns than VTreeMiner and TPMiner. Tatikonda et al. [21] proposed efficient techniques for parallel mining of trees on multicore systems. Since we do not aim at parallel tree mining, their system is not proper for our comparisons.

We used several real-world datasets from different areas to evaluate the efficiency of TPMiner. The datasets do neither have noise nor missing values. Table 1 reports basic statistics of our real-world datasets.

The first real-world dataset is CSLOGS [26] that contains the web access trees of the CS department of the Rensselaer Polytechnic Institute. It has 59,691 trees, 716,263 vertices and 13,209 vertex labels. Each distinct label corresponds to the




Fig. 10: Comparison over CSLOGS32241.




Fig. 11: Comparison over Prions.

URLs of a web page. As discussed in [20], when *per-occurrence* frequency is used, none of the algorithms can find meaningful patterns, because by decreasing *minsup*, suddenly lots of frequent patterns with many occurrences appear in the dataset that makes the mining task practically intractable. The authors of [20] progressively reduced the dataset and generated the CSLOGS32241 dataset that contains 32,241 trees. Figure 10 compares the algorithms over CSLOGS32241. For all given values of *minsup*, VTreeMiner does not terminate within a reasonable time (i.e., 1 day!), therefore, the figure does not contain it. TPMiner is faster than MB3Miner by a factor of 5-20 and it requires significantly less memory cells. In order to have a comparison with VTreeMiner, we tested the algorithms for *minsup* = 1000; while TPMiner finds frequent patterns within around 1.3 seconds, VTreeMiner takes more than 1000 seconds to find the same patterns.

The second real-world dataset used in this paper is Prions that describes a protein ontology database for Human Prion proteins in XML format [19]. The authors of [20] converted it into a tree-structured dataset by considering tags as vertex labels. It has 17,551 wide trees. Figure 11 reports the empirical results over this dataset, where TPMiner is faster than VTreeMiner by a factor of 5-5.2, and it is faster than MB3Miner by a factor of 7.3-11. The third real-world dataset is a dataset of IP mul-




Fig. 12: Comparison over NASA.

ticast. The NASA dataset consists of MBONE multicast data that was measured during the NASA shuttle launch between the 14th and 21st of February, 1999 [4] and [5]. It has 333 distinct vertex labels where each vertex label is an IP address. The data was sampled from this NASA dataset with 10 minutes sampling interval and has 1,000 trees. In this dataset, large frequent patterns are found at high minimum support values. As depicted in Figure 12, over this dataset, TPCMiner is 3-4 times faster than VTreeMiner and both methods are significantly faster than MB3Miner. At  $minsup = 902$ , MB3Miner fails.

We also evaluated the efficiency of the proposed algorithm on several synthetic datasets generated by the method described in [26]. The synthetic data generation program mimics the web site browsing behavior of the user. First a master web site browsing tree is built and then the subtrees of the master tree are generated. The program is adjusted by 5 parameters: (i) the number of labels ( $N$ ), (ii) the number of vertices in the master tree ( $M$ ), (iii) the maximum fan-out of a vertex in the master tree ( $F$ ), (iv) the maximum depth of the master tree ( $D$ ), and (v) the total number of trees in the dataset ( $T$ ). Figure 13 compares the algorithms over the synthetic datasets. The first synthetic dataset is D10 and uses the following default values for the parameters:  $N = 100$ ,  $M = 10,000$ ,  $D = 10$ ,  $F = 10$  and  $T = 100,000$ . Over this dataset, TPCMiner is around 3 times faster than VTreeMiner and VTreeMiner is slightly faster than MB3Miner. The next synthetic dataset is D5, where  $D$  is set to 5 and for the other parameters, the default values are used. Over this dataset, at  $minsup = 10, 20$  and  $30$ , MB3Miner is aborted due to lack of memory. We also evaluated the effect of  $M$ . We set  $M$  to 100,000 and used the default values for the other parameters and generated the M100k dataset. Over this dataset, TPCMiner is faster than MB3Miner by a factor of 2-3, and both TPCMiner and MB3Miner are significantly faster than VTreeMiner.

*Discussion.* Our extensive experiments report that TPCMiner always outperforms well-known existing algorithms. Furthermore, there are several cases where TPCMiner by order of magnitude is more efficient than any specific given algorithm. TPCMiner and VTreeMiner require significantly less memory cells than MB3Miner. This is due to the different large data-structures used by MB3Miner such as the so-called EL, OC and VOL data structures and also to the breadth-first search (BFS) strategy followed

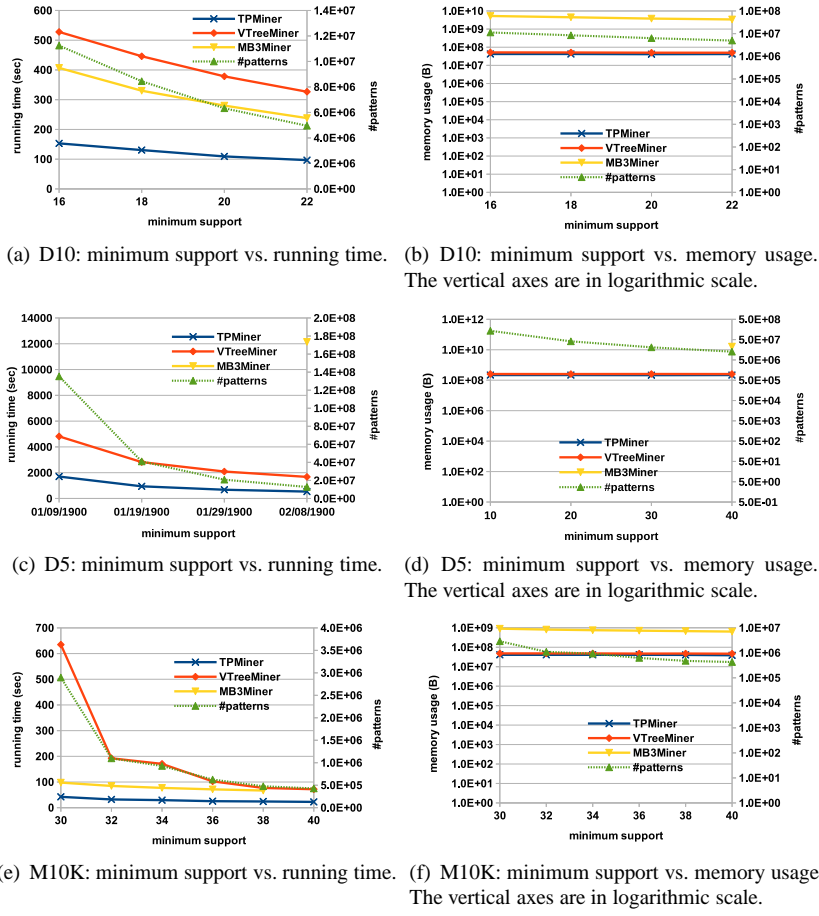


Fig. 13: Comparison over synthetic datasets.

by MB3Miner [20]. Although TPMiner uses a more compact representation of occurrences than VTreeMiner, this is hardly noticeable in the charts. The reason is that the memory use is dominated by the storage of the frequent patterns.

In our experiments, we can distinguish two cases. First, over datasets such as NASA and D5 (in particular for low values of *minsup*), the Tree Model Guided technique used by MB3Miner does not significantly reduce the state space, therefore, the algorithm fails or it does not terminate within a reasonable time. In such cases, TPMiner find all frequent patterns very effectively. Second, over very large datasets (such as CSLOGS32241) or dense datasets (such as M100K) where patterns have many occurrences, TPMiner becomes faster than VTreeMiner by order of magnitude. This is due to the ability of TPMiner in frequency counting of patterns with many occurrences. As discussed earlier, the *occ* data-structure used by TPMiner can

often represent and handle exponentially many occurrences with a single **occ** element, while in VTreeMiner these occurrences are represented and handled one by one.

## 7 Conclusion

In this paper, we proposed an efficient algorithm for subtree homeomorphism with application to frequent pattern mining. We developed a compact data-structure, called **occ**, that effectively represents/encodes several occurrences of a tree pattern. We then defined efficient join operations on **occ** that help us to count occurrences of tree patterns according to occurrences of their proper subtrees. Based on the proposed subtree homeomorphism method, we introduced TPMiner, an effective algorithm for finding frequent tree patterns. We evaluated the efficiency of TPMiner on several real-world and synthetic datasets. Our extensive experiments show that TPMiner always outperforms well-known existing algorithms, and there are several situations where the improvement compared to existing algorithms is significant.

**Acknowledgements** We are grateful to Professor Mohammed Javeed Zaki for providing the VTreeMiner code, the CSLOGS datasets and the TreeGenerator program, to Dr Henry Tan for providing the MB3Miner code, to Dr Fedja Hadzic for providing the Prions dataset and to Professor Jun-Hong Cui for providing the NASA dataset. Finally, we would like to thank Dr Morteza Haghir Chehreghani for his discussion and suggestions.

## References

1. Asai, T., Abe, K., Kawasoe, S., Arimura, H., Satamoto, H., Arikawa, S.: Efficient substructure discovery from large semi-structured data. In: Proceedings of the Second SIAM International Conference on Data Mining (SDM), pp. 158–174. SIAM (2002)
2. Balcazar, J.L., Bifet, A., Lozano, A.: Mining frequent closed rooted trees. *Machine Learning* **78**(1-2), 1–33 (2010)
3. Bille, P., Gørtz, I.: The tree inclusion problem: in linear space and faster. *ACM Transactions on Algorithms* **7**(3), 1–47 (2011)
4. Chalmers, R., Almeroth, K.: Modeling the branching characteristics and efficiency gains of global multicast trees. In: Proceedings of the 20th IEEE International Conference on Computer Communications (INFOCOM), pp. 449–458 (2001)
5. Chalmers, R.C., Member, S., Almeroth, K.C.: On the topology of multicast trees. *IEEE/ACM Transactions on Networking* **11**, 153–165 (2003)
6. Chaoji, V., Hasan, M.A., Salem, S., Zaki, M.J.: An integrated, generic approach to pattern mining: data mining template library. *Data Mining and Knowledge Discovery* **17**(3), 457–495 (2008)
7. Chehreghani, M.H.: Efficiently mining unordered trees. In: Proceedings of the 11th IEEE International Conference on Data Mining (ICDM), pp. 111–120 (2011)
8. Chehreghani, M.H., Chehreghani, M.H., Lucas, C., Rahgozar, M.: OInduced: an efficient algorithm for mining induced patterns from rooted ordered trees. *IEEE Transactions on Systems, Man, and Cybernetics, Part A* **41**(5), 1013–1025 (2011)
9. Chi, Y., Muntz, R.R., Nijssen, S., Kok, J.N.: Frequent subtree mining - an overview. *Fundamenta Informaticae* **66**(1-2), 161–198 (2005)
10. Chi, Y., Yang, Y., Muntz, R.R.: Indexing and mining free trees. In: Proceedings of the Third IEEE International Conference on Data Mining (ICDM), pp. 509–512 (2003)
11. Cui, J., Kim, J., Maggiorini, D., Boussetta, K., Gerla, M.: Aggregated multicast - a comparative study. In: Proceedings of the Second International IFIP-TC6 Networking Conference on Networking Technologies, Services, and Protocols; Performance of Computer and Communication Networks; and Mobile and Wireless Communications (NETWORKING), pp. 1032–1044 (2002)

12. Diestel, R.: Graph Theory, 4th ed. Springer-Verlag, Heidelberg (2010)
13. Dietz, P.F.: Maintaining order in a linked list. In: Proceedings of the 14th ACM Symposium on Theory of Computing (STOC), pp. 122–127 (1982)
14. Ivancsy, R., Vajk, I.: Frequent pattern mining in web log data. *Acta Polytechnica Hungarica* **3**(1), 77–90 (2006)
15. Kilpelainen, P., Mannila, H.: Ordered and unordered tree inclusion. *SIAM Journal of Computing* **24**(2), 340–356 (1995)
16. Miyahara, T., Suzuki, Y., Shoudai, T., Uchida, T., Takahashi, K., Ueda, H.: Discovery of maximally frequent tag tree patterns with contractible variables from semistructured documents. In: Proceedings of the 8th Pacific Asia Conference on Knowledge Discovery and Data Mining (PAKDD), pp. 133–144 (2004)
17. Nijssen, S., Kok, J.N.: Efficient discovery of frequent unordered trees. In: Proceedings of the First International Workshop on Mining Graphs, Trees, and Sequences (MGTS), pp. 55–64 (2003)
18. Qin, L., Yu, J.X., Ding, B.: TwigList: make twig pattern matching fast. In: Proceedings of the 12th International Conference on Database Systems for Advanced Applications (DASFAA), pp. 850–862 (2007)
19. Sidhu, A.S., Dillon, T.S., Chang, E.: Protein ontology. In: Z. Ma, J.Y. Chen (eds.) *Database Modeling in Biology: Practices and Challenges*, pp. 39–60. Springer-Verlag (2006)
20. Tan, H., Hadzic, F., Dillon, T.S., Chang, E., Feng, L.: Tree model guided candidate generation for mining frequent subtrees from XML documents. *ACM Transaction on Knowledge Discovery from Data (TKDD)* **2**(2), 43 pages (2008). DOI 10.1145/1376815.1376818
21. Tatikonda, S., Parthasarathy, S.: Mining tree-structured data on multicore systems. *The Proceedings of the VLDB Endowment (PVLDB)* **2**(1), 694–705 (2009)
22. Tatikonda, S., Parthasarathy, S., Kurc, T.M.: TRIPS and TIDES: new algorithms for tree mining. In: Proceedings of the 15th ACM International Conference on Information and Knowledge Management (CIKM), pp. 455–464 (2006)
23. Wang, C., Hong, M., Pei, J., Zhou, H., Wang, W., Shi, B.: Efficient pattern-growth methods for frequent tree pattern mining. In: Proceedings of the 8th Pacific Asia Conference on Knowledge Discovery and Data Mining (PAKDD), pp. 441–451 (2004)
24. Xiao, Y., Yao, J.F., Li, Z., Dunham, M.H.: Efficient data mining for maximal frequent subtrees. In: Proceedings of the Third IEEE International Conference on Data Mining (ICDM), pp. 379–386 (2003)
25. Zaki, M.J.: Efficiently mining frequent embedded unordered trees. *Fundamenta Informaticae* **66**(1-2), 33–52 (2005)
26. Zaki, M.J.: Efficiently mining frequent trees in a forest: algorithms and applications. *IEEE Transaction on Knowledge and Data Engineering* **17**(8), 1021–1035 (2005)
27. Zaki, M.J., Aggarwal, C.C.: XRules: an effective algorithm for structural classification of XML data. *Machine Learning* **62**(1-2), 137–170 (2006)